

CLAIMS

What is claimed is:

1. A processor-implemented method for determining causal relations between a plurality of intercommunicating nodes, comprising:
 - inputting trace data that describe inter-node communication, the trace data including for each message sent between nodes a timestamp that indicates a time at which the message was sent, a source identifier that identifies a node from which the message was sent, and a destination identifier that identifies a node to which the message was sent;
 - determining for each of one or more nodes whether one or more causal relations exist between a first set of messages destined to the node and a second set of messages sourced from the node and destined to at least one other node, wherein a causal relation exists as a function of a probability distribution of delay values that are differences between timestamps of messages in the second set and timestamps of messages in the first set; and
 - generating a processor-readable representation of the nodes and causal relations.
2. The method of claim 1, wherein the nodes are represented as vertices and the causal relations are represented as edges that connect the vertices, the method further comprising associating with edges information that indicates estimated time delays between nodes as a function of the delay values.
3. The method of claim 2, wherein a causal relation exists if timestamps of a selected number of messages in the second set are offset from timestamps of messages in the first set by approximately equal delay values.
4. The method of claim 3, further comprising, if a causal relation is determined between messages to a first node and messages from the first node to a second node and associated, approximately equal delay values are negative, then discarding the causal relation.
5. The method of claim 3, further comprising:

determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;

determining from the timestamps in the trace data a second indicator function of times of the messages sourced from the source node and destined to the destination node;

- 5 determining a cross-correlation function of the first and second indicator functions;
determining whether a relative peak is present in the cross-correlation function;
and

if a relative peak is present, associating the delay value at the relative peak with an edge from a vertex representing the source node to a vertex representing the destination
10 node.

6. The method of claim 5, further comprising if a relative peak is present, associating a weight value with the edge from the vertex that represents the source node to the vertex that represents the destination node, wherein the weight value is a function of a number of
15 pairs of messages in the first set and second set having approximately equal delay values.

7. The method of claim 5, further comprising if a relative peak is present, associating a weight value with the edge from the vertex that represents the source node to the vertex that represents the destination node, wherein the weight value is a function of
20 characteristics of the relative peak.

8. The method of claim 3, further comprising:
determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;
25 convolving the first indicator function with a distribution function that models an expected variance of the approximately equal delay values, whereby a convolved-first indicator function is generated;
determining from the timestamps in the trace data a second indicator function of times of the messages sourced from the source node and destined to the destination node;
30 convolving the second indicator function with the distribution function, whereby a convolved-second indicator function is generated;
determining a cross-correlation function of the convolved-first and convolved-second indicator functions;

determining whether a relative peak is present in the cross-correlation function;
and

if a relative peak is present, associating the delay value at the relative peak with an edge from a vertex representing the source node to a vertex representing the destination node.

9. The method of claim 1, further comprising, if a first causal relation is determined between messages to a first node and messages from the first node to a second node, a second causal relation is determined between messages to the second node and messages from the second node to the first node, and if some other causal relation is detected between messages from the first node to the second node and messages from the second node to a node other than the first node, then discarding the second causal relation.

10. The method of claim 1, wherein the nodes are represented as vertices and the causal relations are represented as edges that connect the vertices, the method further comprising:

associating with edges information that indicates estimated time delays between nodes as a function of the delay values; and

limiting to a selected maximum number, vertices that represent a particular node in a path from an initial vertex.

11. The method of claim 1, wherein determining causal relations comprises:

for messages from a source node and destined to one or more destination nodes, forming one or more subsets of messages, wherein messages in each subset are destined to a common destination node;

determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;

determining from the timestamps in the trace data one or more second indicator functions, each second indicator function corresponding to one of the subsets of messages and being a function of times of the messages sourced from the source node and destined to the common destination node of messages in the subset;

determining one or more cross-correlation functions from the first indicator function and the one or more second indicator functions;

determining in each cross-correlation function whether a relative peak is present;
and
generating a representation of a causal relation for each relative peak.

- 5 12. The method of claim 1, wherein determining causal relations comprises:
for each second set of messages from a source node and destined to one or more
destination nodes, forming one or more subsets of messages, wherein messages in each
subset are destined to a common destination node;
determining from the timestamps in the trace data a first indicator function of times
10 of the messages destined to the source node;
determining from the timestamps in the trace data one or more second indicator
functions, each second indicator function corresponding to one of the subsets of messages
and being a function of times of the messages sourced from the source node and destined
to the common destination node of messages in the subset;
15 determining one or more cross-correlation functions from the first indicator
function and the one or more second indicator functions;
determining in each cross-correlation function whether a relative peak is present;
generating a representation of a causal relation for each relative peak;
if a relative peak is present in a cross-correlation function of the first indicator
20 function and one of the one or more second indicator functions, removing from the first set
messages associated with the relative peak, and removing messages associated with the
relative peak from the subset that corresponds to the one of the one or more second
indicator functions; and
repeating for each subset of messages the steps of determining a cross-correlation
25 function, determining a relative peak, representing each relative peak, and removing
messages until a relative peak is not present.

13. The method of claim 12, wherein the repeating step is performed while at least a
selected minimum number of messages remain in the first set of messages and in the
30 subset of messages.

14. The method of claim 1, wherein the nodes are represented as vertices and the causal relations are represented as edges that connect the vertices, and the step of determining causal relations further comprises:

- for each second set of messages from a source node and destined to one or more destination nodes, forming one or more subsets of messages, wherein messages in each subset are destined to a common destination node;
- determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;
- determining from the timestamps in the trace data one or more second indicator functions, each second indicator function corresponding to one of the subsets of messages and being a function of times of the messages sourced from the source node and destined to the common destination node of messages in the subset;
- determining one or more cross-correlation functions from the first indicator function and the one or more second indicator functions;
- determining in each cross-correlation function whether a relative peak is present;
- if a relative peak is present in the cross-correlation function of the first indicator function and one of the one or more second indicator functions, adding a delay value at the relative peak to a set of delay values associated with an edge from a vertex representing the source node to a vertex representing the destination node of the subset of messages that corresponds to the one of the one or more second indicator functions, removing from the first set messages associated with the relative peak, and removing from the subset messages associated with the relative peak; and
- repeating for each subset of messages the steps of determining a cross-correlation function, determining a relative peak, representing each relative peak, and removing messages until a relative peak is not present.

15. The method of claim 1, wherein determining causal relations comprises:

- for messages from a source node and destined to one or more destination nodes, selecting a subset of the one or more destination nodes, wherein each destination node in the subset is a destination in at least a selected number of the messages, and forming one or more subsets of messages, wherein messages in each subset are destined to a common destination node in the selected subset of destination nodes;

determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;

determining from the timestamps in the trace data one or more second indicator functions, each second indicator function corresponding to one of the subsets of messages
5 and being a function of times of the messages sourced from the source node and destined to the common destination node of messages in the subset;

determining one or more cross-correlation functions from the first indicator function and the one or more second indicator functions;

determining in each cross-correlation function whether a relative peak is present;
10 and

generating a representation of a causal relation for each relative peak.

16. An apparatus for determining causal relations between a plurality of intercommunicating nodes, comprising:

15 means for inputting trace data that describe inter-node communication, the trace data including for each message sent between nodes a time value that indicates a time at which the message was sent, a source identifier that identifies a node from which the message was sent, and a destination identifier that identifies a node to which the message was sent;

20 means for determining for each of one or more nodes whether one or more causal relations exist between a first set of messages destined to the node and a second set of messages sourced from the node and destined to at least one other node, wherein a causal relation exists as a function of a probability distribution of delay values that are differences between timestamps of messages in the second set and timestamps of
25 messages in the first set; and

means for generating a processor-readable representation of the nodes and causal relations.

17. The apparatus of claim 16, wherein a causal relation exists if timestamps of a
30 selected number of messages in the second set are offset from timestamps of messages in the first set by approximately equal delay values, and the nodes are represented as vertices and the causal relations are represented as edges that connect the vertices, the apparatus

further comprising means for associating with edges information that indicates estimated time delays between nodes as a function of the delay values.

18. The apparatus of claim 17, further comprising means for weighting a
5 representation of a causal relation as a function of a detected number of messages in the second set that are offset from timestamps of the detected number of messages in the first set by approximately equal delay values.

19. The apparatus of claim 17, further comprising means for eliminating representation
10 of a cycle of vertices and edges.

20. An article of manufacture for determining causal relations between a plurality of intercommunicating nodes, comprising:
a computer-readable medium configured with instructions for causing a processor-
15 based system to perform the steps of,

inputting trace data that describe inter-node communication; the trace data including for each message sent between nodes a timestamp that indicates a time at which the message was sent, a source identifier that identifies a node from which the message was sent, and a destination identifier that identifies a node to which
20 the message was sent;

determining for each of one or more nodes whether one or more causal relations exist between a first set of messages destined to the node and a second set of messages sourced from the node and destined to at least one other node, wherein a causal relation exists as a function of a probability distribution of delay values
25 that are differences between timestamps of messages in the second set and timestamps of the selected number of messages in the first set; and

generating a processor-readable representation of the nodes and causal relations.

30 21. The article of manufacture of claim 20, wherein the nodes are represented as vertices and the causal relations are represented as edges that connect the vertices, and the computer-readable medium is further configured with instructions for causing a processor-

based system to perform the step of associating with edges information that indicates estimated time delays between nodes as a function of the delay values.

22. The article of manufacture of claim 21, wherein a causal relation exists if
5 timestamps of a selected number of messages in the second set are offset from timestamps of messages in the first set by approximately equal delay values.

23. The article of manufacture of claim 22, wherein the computer-readable medium is further configured with instructions for causing a processor-based system to perform the
10 steps of:

determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;

determining from the timestamps in the trace data a second indicator function of times of the messages sourced from the source node and destined to the destination node;

15 determining a cross-correlation function of the first and second indicator functions;

determining whether a relative peak is present in the cross-correlation function;

and

if a relative peak is present, associating the delay value at the relative peak with an edge from a vertex representing the source node to a vertex representing the destination

20 node.

24. The article of manufacture of claim 23, wherein the computer-readable medium is further configured with instructions for causing a processor-based system to perform the step of associating, if a relative peak is present, a weight value with the edge from the
25 vertex that represents the source node to the vertex that represents the destination node, wherein the weight value is a function of a number of pairs of messages in the first set and second set having approximately equal delay values.

25. The article of manufacture of claim 23, wherein the computer-readable medium is further configured with instructions for causing a processor-based system to perform the
30 step of associating a weight value with the edge from the vertex that represents the source node to the vertex that represents the destination node if a relative peak is present, wherein the weight value is a function of characteristics of the relative peak.

26. The article of manufacture of claim 21, wherein the computer-readable medium is further configured with instructions for causing a processor-based system to perform the step of discarding a causal relation if the causal relation is determined between messages to a first node and messages from the first node to a second node and associated, approximately equal delay values are negative.

27. The article of manufacture of claim 21, wherein the computer-readable medium is further configured with instructions for causing a processor-based system to perform the step of limiting to a selected maximum number, vertices that represent a particular node in a path from an initial vertex.

28. The article of manufacture of claim 20, wherein the computer-readable medium is further configured with instructions for causing a processor-based system to perform the step of discarding a particular causal relation, if a first causal relation is determined between messages to a first node and messages from the first node to a second node, the particular causal relation is determined between messages to the second node and messages from the second node to the first node, and if some other causal relation is detected between messages from the first node to the second node and messages from the second node to a node other than the first node.

29. The article of manufacture of claim 20, wherein the instructions for determining causal relations further comprise instructions for:

for messages from a source node and destined to one or more destination nodes, forming one or more subsets of messages, wherein messages in each subset are destined to a common destination node;

determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;

determining from the timestamps in the trace data one or more second indicator functions, each second indicator function corresponding to one of the subsets of messages and being a function of times of the messages sourced from the source node and destined to the common destination node of messages in the subset;

determining one or more cross-correlation functions from the first indicator function and the one or more second indicator functions;

determining in each cross-correlation function whether a relative peak is present; and

5 generating a representation of a causal relation for each relative peak.

30. The article of manufacture of claim 20, wherein the instructions for determining causal relations further comprise instructions for:

for each second set of messages from a source node and destined to one or more
10 destination nodes, forming one or more subsets of messages, wherein messages in each subset are destined to a common destination node;

determining from the timestamps in the trace data a first indicator function of times of the messages destined to the source node;

determining from the timestamps in the trace data one or more second indicator
15 functions, each second indicator function corresponding to one of the subsets of messages and being a function of times of the messages sourced from the source node and destined to the common destination node of messages in the subset;

determining one or more cross-correlation functions from the first indicator function and the one or more second indicator functions;

20 determining in each cross-correlation function whether a relative peak is present; generating a representation of a causal relation for each relative peak;

if a relative peak is present in a cross-correlation function of the first indicator function and one of the one or more second indicator functions, removing from the first set messages associated with the relative peak, and removing messages associated with the
25 relative peak from the subset that corresponds to the one of the one or more second indicator functions; and

repeating for each subset of messages the steps of determining a cross-correlation function, determining a relative peak, representing each relative peak, and removing messages until a relative peak is not present.

30

31. The article of manufacture of claim 20, wherein the instructions for determining causal relations further comprise instructions for:

- for messages from a source node and destined to one or more destination nodes,
selecting a subset of the one or more destination nodes, wherein each destination node in
the subset is a destination in at least a selected number of the messages, and forming one
or more subsets of messages, wherein messages in each subset are destined to a common
5 destination node in the selected subset of destination nodes;
determining from the timestamps in the trace data a first indicator function of times
of the messages destined to the source node;
determining from the timestamps in the trace data one or more second indicator
functions, each second indicator function corresponding to one of the subsets of messages
10 and being a function of times of the messages sourced from the source node and destined
to the common destination node of messages in the subset;
determining one or more cross-correlation functions from the first indicator
function and the one or more second indicator functions;
determining in each cross-correlation function whether a relative peak is present;
15 and
generating a representation of a causal relation for each relative peak.